UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/715,225	11/17/2003	Arun Kwangil Iyengar	YOR920030488US1 (163-16)	5015

24336 7590 01/17/2008

KEUSEY, TUTUNJIAN & BITETTO, P.C.
20 CROSSWAYS PARK NORTH
SUITE 210
WOODBURY, NY 11797

EXAMINER	
TSAI, SHENG JEN	

ART UNIT	PAPER NUMBER
2186	

MAIL DATE	DELIVERY MODE
01/17/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/715,225	IYENGAR ET AL.	
	Examiner	Art Unit	
	Sheng-Jen Tsai	2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *21 November 2007*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-26* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-26* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *17 November 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. This Office Action is taken in response to Applicants' Amendments and Remarks filed on November 21, 2007 regarding application 10/715,225 filed on November 17, 2003.

2. Claims 1, 10-11 and 17-18 have been amended.

Claims 1-26 are pending for consideration.

3. ***Response to Amendments and Remarks***

Applicants' amendments and remarks have been fully and carefully considered. In response, a new ground of claim analysis based on the combination of a previously relied on reference (Iyengar et al., U.S. Patent Application Publication 2003/0172236) and a newly identified reference (Hiraoka et al., US 4,733,348) has been made. Refer to the corresponding sections of the following claim analysis for details.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

5. Claims 1-6, 10-12, 16-24 and 26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Iyengar et al. (U.S. Patent Application Publication 2003/0172236, hereinafter referred to as Iyengar), and in view of Hiraoka et al. (US 4,733,348, hereinafter referred to as Hiraoka).

It is noted that, in the following claim analysis, those elements recited by the claims are presented using **bold** font.

As to claim 1, Iyengar discloses **in a system comprised of a plurality of storage elements** [figure 1 shows a system comprising a central cache (110), a remote server (104) and a plurality of processors (106-1~106-N) where each of the processor has a copy of cache as storage element, hence a plurality of storage elements], **a method for maintaining objects in the storage elements** [Methods and Systems for Distributed Caching in Presence of Updates and in Accordance with Holding Times (abstract); the corresponding objects in the storage elements are the contents of the caches] **comprising the steps of**:

maintaining information regarding which storage elements are storing particular objects [the central cache maintains local directories 110 which indicate the contents of local caches. A local directory maintains information about what objects may, but do not necessarily have to be, cached in the corresponding local cache. These local directories 110 allow a central cache to update local caches (paragraph 0030); the central cache 102 stores information from at least one remote server 104. The central cache communicates with a plurality of processors 106 which contain local caches 108. The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (paragraph 0026)] **in a consistency coordinator** [the central cache (figure 1, 110) is the corresponding consistency coordinator; The central cache communicates with the one or more local caches and

coordinates updates to the local caches] **which communicates with the storage elements** [the central cache 102 stores information from at least one remote server 104. The central cache communicates with a plurality of processors 106 which contain local caches 108. The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (paragraph 0026)];

responding to a request to update an object [in step 202 (figure 2), a request for an object is issued (column 5, lines 5-15)] **by using maintained information to determine which of the storage elements may store a copy of the object** [The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (paragraph 0026); In step 304, the central cache coordinates cache updates. That is, the central cache updates all objects it has cached which have changed. In addition, the central cache consults its local directories 110 to see which local caches may contain changed objects. Using local directories, the central cache 102 sends appropriate update messages to local caches (paragraphs 0039-0040)];

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object [The central cache communicates with the one or more local caches and coordinates updates to the local caches, including cache replacement (abstract); In step 304, the central cache coordinates cache updates. That is, the central cache updates all objects it has cached

which have changed. In addition, the central cache consults its local directories 110 to see which local caches may contain changed objects. Using local directories, the central cache 102 sends appropriate update messages to local caches (paragraphs 0039-0040); When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (paragraph 0026); it is to be understood that the term "update," as used herein, is meant not only to include changing the value of a data object in a cache but also may include invalidating the data object or performing some other operation on the object. The central cache may communicate remotely with processes running either on the same processing node or on different processing nodes. That way, several applications running on different processing nodes may communicate with the same cache (paragraph 0010)]; **and delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object** [Iyengar: Updates to cached objects may go through the central cache. In order to update a cache object, the central cache may communicate with the local caches to make sure that all copies are invalidated or updated (paragraph 0011); thus the central cache delays the updating of an object until after it communicates communicate with the local caches to make sure that all copies are invalidated or updated; further, Hiraoka also teaches this aspect: A purge request source processor commonly supplies a purge request signal to other processors so as to cause them to perform TLB purge operations. A purge end signal sent back from other processors is stored in flip-flops in the source processor in units of processors.

The source processor detects the end of TLB purge operations of all processors, in accordance with the statuses of the flip-flops (abstract)] **or (ii) been deemed unresponsible** [taught by Hiraoka, see below].

Regarding claim 1, Iyengar does not explicitly teach delaying an updating of the object until (ii) storage elements are deemed unresponsible.

However, it is well known in the art that a time-out mechanism is commonly used to avoid an infinite waiting period caused by a non-responsive element whose acknowledgement would never arrive. Without such a time-out mechanism, any non-responsive element would prevent the ensuing operations from being executed and cause the system to hang up.

Further, Hiraoka teaches in the invention "Virtual-Memory Multiprocessor System for Parallel Purge Operation" a scheme to send a purge command to a plurality of processors [as shown figure 3] to remove/invalidate a page from a Translation Lookaside Buffer (TLB) [abstract]. Specifically, with respect to claim 1, Hiraoka teaches **delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object** [A purge request source processor commonly supplies a purge request signal to other processors so as to cause them to perform TLB purge operations. A purge end signal sent back from other processors is stored in flip-flops in the source processor in units of processors. The source processor detects the end of TLB purge operations of all processors, in accordance with the statuses of the flip-flops (abstract); thus the acknowledgement from each

storage element is stored and recorded in a corresponding flip-flop, and The source processor delays the updating until after it detects the end of TLB purge operations of all processors, in accordance with the statuses of the flip-flops] **or (ii) been deemed unresponsible** [The above operation can be performed when all the processors 20_0 through 20_3 are present. However, when the processor 20_3 is not present, the following operation is performed. The signal 48_3 representing that the processor 20_3 is not present is set at logic "1". The signal 48_3 of logic "1" is supplied to the OR gate 42_3. The OR gate 42_3 supplies the dummy TLB purge end signal to the AND gate 43. If the processor 20_3 is not present, the processor 20_0 can detect that all the TLB purge operations of the processors 20_0 through 20_2 are completed (column 4, lines 41-50). Note that processor 20_3 is the non-responsible element while processors 20_0 through 20_2 are responsive elements].

Therefore it would have been obvious for persons of ordinary skills in the art at the time of the applicant's invention to also take into consideration the situations where elements may be non-responsible, as demonstrated by Hiraoka, and to incorporate it into the existing scheme disclosed by Iyengar, in order to avoid the incidents that non-responsive elements would prevent the execution of the ensuing operations and cause a system to hang up.

As to claim 2, Iyengar teaches that **the step of maintaining information includes maintaining information regarding which storage elements are storing particular objects in the consistency coordinator** [The central cache contains information about what is stored in local caches 108. When cached data changes, the

central cache 102 is notified. The central cache is then responsible for updating local caches 108 (paragraph 0026); In addition, the central cache consults its local directories 110 to see which local caches may contain changed objects. Using local directories, the central cache 102 sends appropriate update messages to local caches (paragraph 0040)].

As to claim 3, Iyengar teaches that **the consistency coordinator includes multiple nodes** [figure 1 shows the central cache (102), which is by itself one node, is connected to a remote server (104), which serves as another node to facilitate consistency coordination with remote storage elements; The central cache may communicate remotely with processes running either on the same processing node or on different processing nodes. That way, several applications running on different processing nodes may communicate with the same cache (paragraph 0010)] **and each node of the consistency coordinator stores information for a different set of objects** [since local caches require extra space and may thus in some situations be of limited size, it is preferred to have one or more methods for determining which objects to store in a local cache. Such methods, referred to as cache replacement policies, are described below in accordance with the present invention (paragraph 0029)].

As to claim 4, Iyengar teaches that **the storage elements include at least one cache** [figure 1 shows a plurality of nodes of processors (106-1~106-N) where each of the processor has a copy of cache as storage element].

As to claim 5, Iyengar teaches that **the storage elements are included in a distributed system** [figure 1 shows the configuration of a distributed system; in one

aspect, a distributed caching technique of the invention comprises the use of a central cache and one or more local caches (paragraph 0010)].

As to claim 6, Iyengar teaches **the method as recited in claim 1, further comprising the step of obtaining a lock on the object to be updated before performing the update** [figure 5, steps 502, 504, 506 and 508; with respect to the locking or holding time issue, in another aspect, the invention provides techniques for adaptively determining such time values (paragraph 0013)].

As to claim 10, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 11, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 12, it recites substantially the same limitations as those recited in claim 3, and is rejected by the same reason as applied to claim 3. Refer to "As to claim 3" presented earlier in this section for details.

As to claim 16, it recites substantially the same limitations as those recited in claim 4, and is rejected by the same reason as applied to claim 4. Refer to "As to claim 4" presented earlier in this section for details.

As to claim 17, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 18, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 19, Iyengar et al. teach that **the system as recited in claim 18, further comprising a writer, which updates the object to be updated** [it is to be understood that the term "update," as used herein, is meant not only to include changing the value of a data object in a cache (paragraph 0010). It is noted that changing the value of a data object inherently requires a write operation, hence a writer].

As to claim 20, Iyengar et al. teach that **the writer resides on a same node as a storage element** [for example, the central cache which would update the cache data objects as shown in figure 1].

As to claim 21, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 22, it recites substantially the same limitations as those recited in claim 1, and is rejected by the same reason as applied to claim 1. Refer to "As to claim 1" presented earlier in this section for details.

As to claim 23, Iyengar teaches that **the system as recited in claim 18, further comprising at least one content provider** [for example, the central cache or the remote server as shown in figure 1; for instance, a cache may be implemented as a

server in a network (e.g., a cache server or proxy caching server in a World Wide Web or Internet environment) (paragraph 0009)].

As to claim 24, Iyengar teaches that **the content provider resides on a same node as a storage element** [for example, the central cache which would update the cache data objects as shown in figure 1].

As to claim 26, it recites substantially the same limitations as those recited in claim 4, and is rejected by the same reason as applied to claim 4. Refer to "As to claim 4" presented earlier in this section for details.

6. Claims 7-9, 14-15 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Iyengar et al. (U.S. Patent Application Publication 2003/0172236, hereinafter referred to as Iyengar), in view of Hiraoka et al. (US 4,733,348, hereinafter referred to as Hiraoka), and further in view of Chang et al. (US Patent Application Publication 2005/0128960, hereinafter referred to as Chang).

As to claims 7-9, Iyengar in view of Hiraoka does not mention that **sending heart beat messages to obtain availability information to and from a storage element**.

However, Chang discloses in their invention "Method for Determination of Remote Adapter and/or Node Liveness" a heart beat message protocol for the determination of node liveness in a distributed data processing system [abstract; figures 6-8; paragraph 0017].

Chang teaches that using hear beat messages allows early detections of any failure component and prompt recovery operations to maintain high availability of system [Chang, paragraph 0003].

Therefore it would have been obvious for persons of ordinary skills in the art at the time of the applicant's invention to recognize the benefits using hear beat messages to identify faulty components as soon as possible, as demonstrated by Chang, and to incorporate it into the existing apparatus and method disclosed by Iyengar in view of Hiraoka, to further improve the availability and reliability of the system.

As to claims 14-15, they recite substantially the same limitations as those recited in claims 7-9, and are rejected by the same reason as applied to claims 7-9. Refer to "As to claims 7-9" presented earlier in this section for details.

As to claim 25, it recites substantially the same limitations as those recited in claims 7-9, and is rejected by the same reason as applied to claims 7-9. Refer to "As to claims 7-9" presented earlier in this section for details.

7. ***Related Prior Art of Record***

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Gannon et al., (US 5,265,232), "Coherency Control by Data Invalidation in Selected Processor Caches without Broadcasting to Processor Caches not Having the Data."

- Chang et al., (US 5,398,325), "Method and Apparatus for Improving Cache Consistency Using a Single Copy of a Cache tag Memory in Multi Processor Computer Systems."
- Butts, Jr. et al., (US 5,303,362), "Coupled Memory Multiprocessor Computer System Including Cache Coherency Management Protocols."
- Hayes et al., (US 6,073,212), "Reducing Bandwidth and Areas Needed for Non-Inclusive Memory Hierarchy by Using Dual Tags."
- McDonald et al., (US 6,012,127), "Multiprocessor Computing Apparatus with Optional Coherency Directory."
- Teramotop, (US 6,848,023), "Cache Directory Configuration Method and Information Processing Device."

Conclusion

8. Claims 1-26 are rejected as explained above.

9. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

10. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

January 16, 2008

MATTHEW KIM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100